

NP-HARDNESS OF IDEAL LATTICE PROBLEMS

DANIEL E. MARTIN

ABSTRACT. We establish the worst-case hardness of several ideal lattice problems (including **SVP** and **CVP**) in the ℓ_2 norm by providing a dimension-preserving, deterministic polynomial time reduction from their generic lattice versions. The reduction constructs an ideal lattice in the canonical embedding of a number field that approximates some input lattice up to scaling and orthogonal transformation. The integers defining the ideal and the ambient number ring, in particular its discriminant, are all polynomial in bit length relative to the generic input lattice. Furthermore, the ideal is invertible, the ring is monogenic, and the number field is totally real. If the number ring is also required to be a full ring of integers, the reduction conjecturally succeeds in bounded-error quantum polynomial time.

1. INTRODUCTION

The theoretical complexity of ideal-lattice reduction has come under scrutiny due to its cryptographic implications. In this paper, we prove the **NP**-hardness of two such problems, called **SVP** and **CVP**, under appropriate parameters.

Beginning with the work of Ajtai–Dwork [AD97] and Hoffstein–Pipher–Silverman [HPS98], lattices have been a significant source of hard problems in cryptography. Lattice-based encryption is unique in its benefits—presumed quantum security, full homomorphicity [Gen09; BV11], average-case hardness [Ajt96; Reg05; Pei09]—but natively suffers from large key sizes. As a remedy, Hoffstein thought to work over cyclotomic rings $\mathbb{Z}[x]/(x^n - 1)$, where the public key becomes a single element defined by only n integers rather than n^2 (the number of integers that define a generic lattice). He developed this idea alongside Pipher and Silverman into the NTRU cryptosystem. Following NTRU, the theory of cyclotomic ring-based encryption was advanced notably by Micciancio [Mic07] and Peikert and Rosen [PR06]. The notion was generalized by Micciancio and Lyubashevsky to the rings $\mathbb{Z}[x]/(f(x))$ for monic irreducible polynomials $f(x)$ [LM06; Mic07]. Further development by Lyubashevsky, Peikert, Regev, Stephens-Davidowitz, Rosca, Stehlé, Wallet, and others [LPR10; PRS17; RSW18] brought theoretical cryptography to its modern state: lattices are drawn from the canonical embedding of ideals in any ring of integers, or more generally, from any number ring provided the ideal is invertible. These lattices underlie the theoretical hardness of problems like Ring **LWE** [Reg05; LPR10], Polynomial **LWE** [Ste+09], and Order **LWE** [RSW18; Bol+19].

Date: September 13, 2026.

2020 Mathematics Subject Classification. Primary: 68Q17, 11H06, 11R04, 11Y16, 94A60. Secondary: 11Y40, 11R80, 68Q12.

Key words and phrases. lattice reduction, shortest vector problem, closest vector problem, ideal lattice.

This material is based upon work supported by the National Science Foundation under Grant Nos. 2336000 and 2602045.

While the algebraic structure of ideal lattices provides efficiency gains, it may also open the door to new attacks. Despite hints of vulnerability [Bau+17; CGS14; Cra+16; CDW21; PHS19], cryptosystems that derive security from the worst-case hardness of ideal lattice problems are already implemented. This hardness is poorly understood compared to generic lattices, where, for example, the Shortest and Closest Vector Problems (SVP and CVP) are long known to be NP-hard [Ajt98; Emd81]. We address this disparity with a generic-to-ideal lattice reduction algorithm.

Theorem 1.1. *In deterministic polynomial time, Algorithm 1 approximates (to any precision) a scaled, orthogonal transformation of its input lattice with an ideal lattice of the same dimension. The ideal, defined in polynomial bit length, is coprime to the conductor (and thus invertible) in a monogenic, totally real number ring.*

Remark 1.2. There are two choices of subroutine for Algorithm 1: Subroutine 2 is simpler but randomized, while Subroutine 3 is deterministic and justifies Theorem 1.1. The randomized algorithm repeats until the output ring is a full ring of integers. Each iteration conjecturally succeeds with probability roughly 0.323.

The informal moral of Theorem 1.1 is: *The general tools of algebraic number theory do not make ideal lattices inherently more vulnerable than generic lattices.* Conductor coprimality is essential to this claim. Without it, the general tools of algebraic number theory, like unique factorization or the two-generator property (see Section 2.2), do not apply.

Theorem 1.1 does *not* imply that tools specific to certain fields present no new vulnerability. This is because Algorithm 1 does not allow absolute control over the number field in which the output ideal lives. For example, it cannot target a specific field with a preprocessed unit group. There is, however, some control over the field: The algorithm produces a matrix $A \in \text{Mat}_n(\mathbb{Z})$ whose characteristic polynomial, call it $f(x)$, defines the number field and ring. Alternative characteristic polynomials that produce different fields can be obtained by slightly perturbing A or by replacing it with BAB^T for any integer matrix $B \in \text{AGL}_n(\mathbb{Z})A^{-1}$. The resulting ideal lattice still approximates the input up to orthogonal transformation and scaling. The author has not investigated whether these freedoms can produce additional properties like Galois or a near-orthonormal $\mathbb{Z}$ -basis for the ring. Thus the exact degree to which our approach can close the gap between generic lattices and those loved by cryptography is a subject for future research.

As a consequence of Theorem 1.1, hardness results for generic lattice reduction problems in the ℓ_2 norm now apply in the ideal setting. The approximation factors below were proved for SVP_2 and CVP_2 in generic lattices by Wan [Wan26] (building on Bennett and Peikert [BP22]) and Dinur [Din02].

Corollary 1.3. *Restricted to the canonical embedding of ideals coprime to the conductor in totally real, monogenic number rings, SVP_2 and CVP_2 are NP-hard to approximate within factors of $\sqrt{2}$ and $n^{c/\log \log n}$, respectively.*

After we fix notation and definitions in Section 2, we present an intuitive overview of Algorithm 1 and prove its basic properties in Section 3. The main technical hurdle—producing invertible ideals—is avoided until Section 4, which contains our randomized solution. It is expected to produce ideals in a ring of integers, but this experimentally observed high probability of success is essentially equivalent to an unresolved conjecture of Wang and Yu [WY16]. Our deterministic solution is presented in Section 5. This allows us to prove Theorem 1.1 and Corollary 1.3.

2. BACKGROUND

2.1. Lattice problems. We state the search versions of several problems that are amenable to Algorithm 1. The algorithm could just as easily be used to reduce decision versions from the generic to the ideal setting.

For a lattice Λ , recall that $\lambda_i(\Lambda)$ is the minimal radius of a closed ball (in some fixed ℓ_p norm) around the origin that contains i independent vectors in Λ . We use $|\mathbf{v}|_p$ to denote the ℓ_p norm of $\mathbf{v} \in \mathbb{R}^n$.

Definition 2.1. Given $\gamma \in [1, \infty)$ and a basis for a lattice Λ , the *approximate Shortest Vector Problem*, denoted SVP_p or $\gamma\text{-SVP}_p$, asks for a nonzero lattice vector of length at most $\gamma\lambda_1(\Lambda)$.

Definition 2.2. Given $\gamma \in [1, \infty)$ and a basis for an n -dimensional lattice Λ , the *approximate Shortest Independent Vectors Problem*, denoted SIVP_p or $\gamma\text{-SIVP}_p$, asks for n independent lattice vectors of length at most $\gamma\lambda_n(\Lambda)$.

Definition 2.3. Given $\gamma \in [1, \infty)$ and a basis for a lattice Λ with $\gamma\lambda_1(\Lambda) \leq \lambda_2(\Lambda)$, the *unique Shortest Vector Problem*, denoted uSVP_p or $\gamma\text{-uSVP}_p$, is 1-SVP_p in Λ .

Definition 2.4. Given $\gamma \in [1, \infty)$, a basis for a lattice $\Lambda \subset \mathbb{R}^n$, and a target $\mathbf{x} \in \mathbb{R}^n$, the *approximate Closest Vector Problem*, denoted CVP_p or $\gamma\text{-CVP}_p$, asks for $\mathbf{v} \in \Lambda$ with $|\mathbf{v} - \mathbf{x}|_p \leq \gamma \min_{\mathbf{u} \in \Lambda} |\mathbf{u} - \mathbf{x}|_p$.

See [MG12] for a comprehensive overview of these problems and [Pei16] for a focus on cryptographic applications. In [Ben23], Bennett provides a complete survey of complexity results on SVP prior to 2023. Notable subsequent progress has been made by Hair and Sahai [HS26] and Wan [Wan26].

2.2. Number theory. Let $\sigma_1, \dots, \sigma_r : K \hookrightarrow \mathbb{R}$ be the real embeddings of a number field K , and let $\tau_1, \bar{\tau}_1, \dots, \tau_s, \bar{\tau}_s : K \hookrightarrow \mathbb{C}$ be the complex conjugate pairs of non-real embeddings.

Definition 2.5. The *canonical* (or *Minkowski*) *embedding* of K is $\Sigma : K \hookrightarrow \mathbb{R}^r \times \mathbb{C}^s$ defined by $\Sigma(\alpha) = (\sigma_1(\alpha), \dots, \tau_s(\alpha))$.

In contemporary literature, “ideal lattice” typically refers to the canonical embedding of an ideal. We maintain this convention. (Prior to [LPR10], fixing a $\mathbb{Z}$ -basis for $\mathcal{O}$ —typically $1, \dots, x^{n-1}$ when $\mathcal{O} \simeq \mathbb{Z}[x]/(f(x))$ —and using the coefficient embedding was more popular.) Such a lattice has full rank $n := r + 2s$.

We call K *totally real* if $s = 0$. Ideal lattices in totally real fields may be more vulnerable to attacks that leverage unit approximations because they maximize the unit group rank $r + s - 1$ in degree n .

Definition 2.6. The *ring of integers* (or *maximal order*), denoted $\mathcal{O}_K$, consists of elements in K that are roots of monic polynomials in $\mathbb{Z}[x]$. An *order* is a subring of $\mathcal{O}_K$ that is not contained in a proper subfield of K . An order is *monogenic* if it takes the form $\mathbb{Z}[\lambda]$ for some $\lambda \in \mathcal{O}_K$.

There are two properties of ideals $\mathfrak{a} \subseteq \mathcal{O}_K$ that are especially useful in cryptography: they require at most two generators, meaning $\mathfrak{a} = \alpha\mathcal{O}_K + \beta\mathcal{O}_K$, reducing the n^2 integers needed to define a generic lattice in n dimensions to only $2n$; and when $\mathcal{O}_K$ is monogenic, its elements can be multiplied in $O(n \log n)$ operations with the Fast Fourier Transform. These properties also hold for an ideal $\mathfrak{a}$ in a non-maximal

order $\mathcal{O}$ provided $\mathfrak{a}$ is *invertible*, meaning $\mathfrak{a}\mathfrak{b}$ is principal for some ideal $\mathfrak{b} \subseteq \mathcal{O}$. We will target the stronger condition that $\mathfrak{a}$ is coprime to the conductor of $\mathcal{O}$.

Definition 2.7. The *conductor* of an order $\mathcal{O}$ in K is the ideal

$$\mathfrak{f} = \mathfrak{f}(\mathcal{O}) := \{\alpha \in \mathcal{O} \mid \alpha\mathcal{O}_K \subseteq \mathcal{O}\}.$$

Ideals in $\mathcal{O}$ coprime to $\mathfrak{f}$ retain many other desirable number theoretic properties. For example, they represent elements of the class group (or Picard group), factor uniquely into products of primes, and are preserved by a rank $r + s - 1$ unit group. (Note that $\mathfrak{f}(\mathcal{O}_K) = \mathcal{O}_K$, so all ideals in $\mathcal{O}_K$ are coprime to the conductor and invertible.)

Invertibility is an indispensable hypothesis in hardness proofs for ideal variants of Regev’s Learning with Errors problem [LPR10; RSW18; Bol+19; PP24]. If we ignored invertibility, we could easily reduce any generic n -dimensional lattice problem to the ideal setting in our favorite degree- n number field, whatever that may be. Indeed, if $[K : \mathbb{Q}] = n$, a scaled copy of the input lattice can be approximated by some sublattice $\Lambda \subseteq \Sigma(\mathcal{O}_K)$. Since Λ is an ideal in the order $\mathbb{Z} + \text{cov}(\Lambda)\mathcal{O}_K$, we’re done! Unfortunately, such a reduction is useless because the ideal is almost never invertible. It neither supports the belief that ring variants of LWE are hard (because the underlying ideals in every SVP-to-LWE-style security proof are invertible), nor does it make the aforementioned properties of invertible ideals available to generic lattice reduction algorithms.

Our strategy in the next section is completely different. In contrast to the approach above, it often produces ideals in rings of integers. Even when it doesn’t, the ideal can always be made invertible.

3. APPROXIMATING WITH IDEAL LATTICES

3.1. Intuition and an example. Given an $n \times n$ integer matrix M , the goal is to approximate $M\mathbb{Z}^n$ with a lattice $\Sigma(\mathfrak{a})$ for some invertible ideal $\mathfrak{a}$ in an order $\mathcal{O}$.

We start by finding a symmetric matrix A representing the same lattice $M\mathbb{Z}^n$. Any symmetric basis will do. The simplest way to find one is perhaps computing the Smith Normal Form $S = U^{-1}MV$, where S is diagonal and $U, V \in \text{GL}_n(\mathbb{Z})$. Since

$$USU^T\mathbb{Z}^n = (USV^{-1})(VU^T\mathbb{Z}^n) = M\mathbb{Z}^n,$$

the matrix $A := USU^T$ provides a symmetric basis. As a minimal working example,

$$M = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 \\ 3 & -1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}}_{USV^{-1}}$$

yields the symmetric matrix

$$A = \begin{bmatrix} 1 & 0 \\ 3 & -1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 & 3 \\ 0 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 3 \\ 3 & 11 \end{bmatrix}.$$

An eigenvalue of A , say $\lambda \in \overline{\mathbb{Q}}$, has a corresponding eigenvector $\mathbf{v} \in \mathbb{Z}[\lambda]^n$, found by taking a column of the adjugate of $M - \lambda I_n$. In the example, $\lambda = 6 + \sqrt{34}$ and $\mathbf{v} = (3, 5 + \sqrt{34})^T$.

Assuming the characteristic polynomial of A is irreducible, its remaining eigenvalues and eigenvectors are the algebraic conjugates of λ and $\mathbf{v}$. In particular, if

D is the diagonal matrix with diagonal entries $\Sigma(\lambda)$ and E is the matrix of eigenvectors satisfying $AE = ED$, then the columns of E^T are precisely the canonical embeddings of the entries of $\mathbf{v}$, as in

$$E^T = \begin{bmatrix} 3 & 5 + \sqrt{34} \\ 3 & 5 - \sqrt{34} \end{bmatrix}.$$

Furthermore, the $\mathbb{Z}$ -module generated by the entries of $\mathbf{v}$, call it $\mathfrak{a}$, is an ideal in $\mathbb{Z}[\lambda]$. This equates to saying that each entry of $\lambda\mathbf{v}$ is a $\mathbb{Z}$ -linear combination of the entries of $\mathbf{v}$, which is verified by comparing columns on each side of $AE = ED$. In our example, $\mathfrak{a} = 3\mathbb{Z} + (5 + \sqrt{34})\mathbb{Z}$ is a nonprincipal prime ideal of norm 3 in $\mathbb{Z}[\sqrt{34}]$.

By the spectral theorem for symmetric matrices, the normalized eigenvectors of A form an orthonormal basis for $\mathbb{R}^n$. In other words, if C is the diagonal matrix whose i^{th} diagonal entry is the reciprocal of the length of the i^{th} column of E , then EC is an orthogonal matrix. In the example,

$$C = \begin{bmatrix} (68 + 10\sqrt{34})^{-\frac{1}{2}} & 0 \\ 0 & (68 - 10\sqrt{34})^{-\frac{1}{2}} \end{bmatrix}.$$

The final step scales the vector formed from diagonal entries of CD , call it $\mathbf{c}$, by a sufficiently large real number r and approximates the result with a lattice point $\Sigma(\beta) \in \Sigma(\mathbb{Z}[\lambda])$. Crucially, $\Sigma(\beta)$ need not be the *closest* lattice point to $r\mathbf{c}$; that requirement would destroy our algorithm's running time. Provided r dwarfs the entries of $\mathbf{c}$ in magnitude, the crudest of approximations will do (as in lines 7-12 in Algorithm 1).

Figure 1 illustrates the improvement from $r = 20$ to $r = 200$. The input lattice is the left-side image. In the middle image, $20\mathbf{c} \approx (21.1, 1.1)^T$ is approximated by $\Sigma(11 + 2\sqrt{34}) \approx (22.7, -0.7)^T$, and the resulting transformation can be seen by how it affects the outlined fundamental region. The transformation is much closer to orthogonal in the right-side image, where $200\mathbf{c} \approx (210.5, 10.9)^T$ is approximated by $\Sigma(111 + 17\sqrt{34}) \approx (210.1, 11.9)^T$.




FIGURE 1. $M\mathbb{Z}^2$ (left) and two “approximations” $\Sigma(\beta\mathfrak{a})$ using $r = 20$ and 200 .

The reduction is complete: solving a lattice problem in the ℓ_2 norm in $M\mathbb{Z}^n$ is essentially equivalent to solving it in the canonical embedding of $\beta\mathfrak{a}$, an ideal in the monogenic order $\mathbb{Z}[\lambda]$. To see why, observe that applying the orthogonal transformation $(EC)^{-1}$ and scaling by r does not affect relative vector lengths, so $M\mathbb{Z}^n$ may be replaced by $r(EC)^{-1}M\mathbb{Z}^n = r(EC)^{-1}A\mathbb{Z}^n$. But

$$r(EC)^{-1}A = r(EC)^{-1}(AE)E^{-1}$$

$$\begin{aligned}
&= r(EC)^{-1}(ED)E^{-1} && \text{by definition of } E \text{ and } D \\
&= rC^{-1}DE^{-1} \\
&= rDC^{-1}E^{-1} && \text{since } D \text{ and } C^{-1} \text{ are diagonal} \\
&= rD(EC)^T && \text{since } EC \text{ is orthogonal} \\
(3.1) \quad &= rCDE^T && \text{since } C \text{ is diagonal.}
\end{aligned}$$

Now, $E^T \mathbb{Z}^n$ is $\Sigma(\mathfrak{a})$ by definition of $\mathfrak{a}$, and the diagonal of rCD is approximately $\Sigma(\beta)$. Thus $rCDE^T \mathbb{Z}^n$ approximates $\Sigma(\beta\mathfrak{a})$, as desired.

In the example, $\mathbb{Z}[\lambda]$ turned out to be the ring of integers in a degree- n number field, namely the quadratic field $\mathbb{Q}(\sqrt{34})$. In particular, $\mathfrak{a}$ was invertible and coprime to the trivial conductor. This is not uncommon experimentally (which distinguishes our approach from the reduction mentioned at the end of Section 2.2), but it is also not guaranteed. The primary remaining obstacle is to ensure an invertible ideal in a degree- n number field. We offer two methods, one probabilistic and one deterministic, in Sections 4 and 5.

3.2. The algorithm. Forcing $[\mathbb{Q}(\lambda) : \mathbb{Q}] = n$ and $\mathfrak{a}$ invertible is the black box in line 3. Since lines 7 and 10 below only make sense when $\mathbb{Q}(\lambda)$ has n real embeddings, the reader is asked to take this for granted throughout this section. Proofs can be found under Theorems 4.4 and 5.12.

The input κ controls quality of approximation between the initial generic lattice and the output ideal lattice. The manner in which κ relates solutions to SVP, SIVP, uSVP, and CVP in these lattices is the subject of Section 3.3.

Algorithm 1: Approximate an input lattice up to orthogonal transformation with some lattice $\Sigma(\mathfrak{a})$ of an invertible ideal $\mathfrak{a} \subseteq \mathbb{Z}[\lambda]$.

Input: $M \in \text{Mat}_n(\mathbb{Z})$ (lattice basis) and $\kappa \in (1, \infty)$ (condition number bound)
Output: a $\mathbb{Z}$ -basis for an ideal lattice

- 1 $S, U, V \leftarrow$ Smith Normal Form matrices $\triangleright MV = US$ with S diagonal
- 2 $A \leftarrow USU^T$
- 3 $A, r \leftarrow$ output of Subroutine 2 or 3 $\triangleright$ input A, κ ; finds a better A
- 4 $f(x) \leftarrow \det(xI_n - A)$
- 5 $\lambda \leftarrow$ largest eigenvalue of A
- 6 $(\alpha_1, \dots, \alpha_n)^T \leftarrow$ last column of $\text{adj}(\lambda I_n - A)$
- 7 $c_i \leftarrow i^{\text{th}}$ embedding length of $(\lambda\alpha_1, \dots, \lambda\alpha_n)$ $\triangleright$ in ℓ_2 norm; see Remark 3.1
- 8 $r \leftarrow rn\lambda^{n-1} \max_i |c_i|(\sqrt{\kappa} + 1)/(\sqrt{\kappa} - 1)$
- 9 $V \leftarrow$ matrix with i^{th} column $\Sigma(\lambda^{i-1})$ $\triangleright$ Vandermonde matrix
- 10 $\mathbf{c} \leftarrow (c_1^{-1}, \dots, c_n^{-1})^T$
- 11 $\mathbf{b} \leftarrow$ point in $\mathbb{Z}^n$ with $|rV^{-1}\mathbf{c} - \mathbf{b}|_\infty < 1$ $\triangleright$ see Remark 3.1
- 12 $\beta \leftarrow (1, \lambda, \dots, \lambda^{n-1})\mathbf{b}$
- 13 **while** $1 \notin (\beta, f'(\lambda))$ **do** $\triangleright$ skip if using Subroutine 2 and
- 14 | $\beta \leftarrow \beta + 1$ perhaps even 3; see Remark 3.2
- 15 **return** $\Sigma(\alpha_1\beta), \dots, \Sigma(\alpha_n\beta)$

Remark 3.1. Computing $\mathbf{b}$ in line 11 requires floating point arithmetic. Entries of $\mathbf{b}$ need only be within 1 of the corresponding entries of $rV^{-1}\mathbf{c}$, so if μ is the

maximum magnitude among entries of $\mathbf{c}$ and V^{-1} , it suffices to compute $\mathbf{c}$ and V^{-1} to within $(2nr\mu)^{-1}$ of their true value. The total bit length of such a rational approximation is $O(\log(nr\mu))$, roughly the same as its integer truncation.

Remark 3.2. Lines 13–14 make β coprime to the conductor of $\mathbb{Z}[\lambda]$, which holds automatically when Subroutine 2 is used. Even if Subroutine 3 is used, this **while** loop can be skipped if we only care that the output ideal is invertible (as a principal ideal is invertible whether or not it is coprime to $\mathfrak{f}(\mathbb{Z}[\lambda])$). Also, note that $1 \in (\beta, f'(\lambda))$ is stronger than $1 \in (\beta, \mathfrak{f}(\mathbb{Z}[\lambda]))$, but computing the conductor is hard.

Proofs of correctness for Algorithm 1 must be postponed until after Subroutines 2 and 3. For now, note that all steps outside of perhaps lines 3 and 13–14 can be performed in a polynomial (in the input bit length) number of operations. (Storjohann’s deterministic algorithm computes S , U , and V in line 1 in $O(n^4 \log^2(n\mu))$ bit operations, where μ is the maximum entry magnitude of M . The resulting matrices U and V have entry magnitudes $(n\mu)^{O(n)}$ [Sto00, Section 8.1].)

3.3. Approximation quality. The role of κ is to bound the condition number of a matrix that transforms the input lattice into the output lattice. Throughout this section, p defines the norm.

Definition 3.3. The *operator norm* of $T \in \text{GL}_n(\mathbb{R})$ in the ℓ_p norm is

$$\|T\|_p = \max_{|\mathbf{u}|_p=1} |T\mathbf{u}|_p,$$

and the *condition number* is

$$\kappa_p(T) = \|T\|_p \|T^{-1}\|_p = \max_{\mathbf{u}, \mathbf{v} \neq 0} \frac{|T\mathbf{u}|_p / |\mathbf{u}|_p}{|T\mathbf{v}|_p / |\mathbf{v}|_p}.$$

The condition number controls the change in what “short” means in the input versus output lattice. We state the propositions below only for the ℓ_2 norm, although they hold more generally mutatis mutandis. Both are certainly well-known.

Proposition 3.4. *If $A\mathbb{Z}^n$ and $B\mathbb{Z}^n$ are full-rank lattices, then multiplying by AB^{-1} transforms a solution in $B\mathbb{Z}^n$ to SVP_p , SIVP_p , uSVP_p , or CVP_p for the parameter γ into a solution in $A\mathbb{Z}^n$ for $\kappa\gamma$, where $\kappa = \kappa_p(AB^{-1})$.*

Proof. Let $T = AB^{-1}$ and $\kappa = \kappa_p(T)$. Consider $\gamma\text{-SVP}_p$. Let $T\mathbf{v}$ be a shortest vector in $A\mathbb{Z}^n$ and suppose $\mathbf{u}$ solves $\gamma\text{-SVP}_p$ in $B\mathbb{Z}^n$, implying $|\mathbf{u}|_p \leq \gamma|\mathbf{v}|_p$. Thus

$$(3.2) \quad \frac{|T\mathbf{u}|_p}{|T\mathbf{v}|_p} \leq \frac{\gamma|T\mathbf{u}|_p / |\mathbf{u}|_p}{|T\mathbf{v}|_p / |\mathbf{v}|_p} \leq \kappa\gamma,$$

meaning $T\mathbf{u}$ solves $\kappa\gamma\text{-SVP}_p$ in $A\mathbb{Z}^n$.

The exact same argument applies to the other lattice problems. $\square$

Since $\kappa = 1$ cannot be achieved regardless of how large r is in Algorithm 1, Proposition 3.4 suggests we cannot perfectly preserve the error tolerance γ in lattice problems. That is not necessarily the case.

Proposition 3.5. *Suppose $p \in \mathbb{Z} \cup \{\infty\}$, Λ is a rank n integer lattice, and $\gamma^q = \frac{s}{t}$ for positive integers q, s, t with $p \mid q$ if $p < \infty$. In the case of CVP_p , suppose further*

that the target $\mathbf{x}$ satisfies $d\mathbf{x} \in \mathbb{Z}$ with $0 \neq d \in \mathbb{Z}$. If

$$\kappa^q < \begin{cases} 1 + (s\lambda_1(\Lambda)^q)^{-1} & \text{SVP}_p, \text{uSVP}_p \\ 1 + (s\lambda_n(\Lambda)^q)^{-1} & \text{SIVP}_p \\ 1 + (s(dn\lambda_n(\Lambda))^q)^{-1} & \text{CVP}_p, \end{cases}$$

then a lattice problem solution in Λ for the parameter $\kappa\gamma$ is also a solution for γ .

Proof. Let $\lambda_i = \lambda_i(\Lambda)$. Suppose $\kappa^q < 1 + (sn\lambda_1^q)^{-1}$ and that $\mathbf{v} \in \Lambda$ solves $\kappa\gamma$ -SVP_p. Then

$$|\mathbf{v}|_p^q \leq (\kappa\gamma\lambda_1)^q < \left(1 + \frac{1}{s\lambda_1^q}\right) \frac{s\lambda_1^q}{t} \leq \frac{s\lambda_1^q + 1}{t}.$$

Since λ_1^q and $|\mathbf{v}|_p^q$ are integers (recall $\Lambda \subseteq \mathbb{Z}^n$) and there are no integers strictly between $\frac{1}{t}(s\lambda_1^q + 1)$ and $\frac{1}{t}(s\lambda_1^q) = (\gamma\lambda_1)^q$, we conclude that $\mathbf{v}$ solves γ -SVP_p in Λ .

The exact same argument applies to uSVP_p and SIVP_p.

For CVP_p with target $\mathbf{x}$ satisfying $d\mathbf{x} \in \mathbb{Z}^n$, if $\mathbf{v}$ solves $\kappa\gamma$ -CVP_p, the argument above applies beginning with the integer $|d(\mathbf{v} - \mathbf{x})|_p^q$. Note that Λ has covering radius at most $n\lambda_n$. $\square$

Lemma 3.6. For any $A, B \in \text{GL}_n(\mathbb{R})$, if

$$\|A - B\|_p \leq \frac{\kappa - 1}{\|A^{-1}\|_p(\kappa + 1)}$$

for some $\kappa \in \mathbb{R}$, then $\kappa_p(AB^{-1}) \leq \kappa$.

Proof. Let $T = BA^{-1}$. Observe that

$$|T\mathbf{v} - \mathbf{v}|_p \leq \|T - I_n\|_p |\mathbf{v}|_p \leq \|A^{-1}\|_p \|A - B\|_p |\mathbf{v}|_p \leq \frac{\kappa - 1}{\kappa + 1} |\mathbf{v}|_p,$$

where the final inequality is the lemma's hypothesis. We may thus apply the triangle inequality to bound $|T\mathbf{v}|_2$ from above and below:

$$\left(1 - \frac{\kappa - 1}{\kappa + 1}\right) |\mathbf{v}|_p \leq |\mathbf{v}|_p - |T\mathbf{v} - \mathbf{v}|_p \leq |T\mathbf{v}|_p \leq |\mathbf{v}|_p + |T\mathbf{v} - \mathbf{v}|_p \leq \left(1 + \frac{\kappa - 1}{\kappa + 1}\right) |\mathbf{v}|_p$$

The ratio of the right-side factor to the left-side factor is exactly κ . Inverting T does not change its condition number, so $\kappa_p(AB^{-1}) \leq \kappa$. $\square$

4. RANDOMIZED SUBROUTINE TO FIND A RING OF INTEGERS

If A is perturbed randomly, we expect its characteristic polynomial, call it $f(x)$, to be irreducible with probability approaching 1 as the dimension n grows. This was confirmed by Eberhard for generic matrices [Ebe22] and Ferber, Jain, Sah, and Sawhney for symmetric matrices [Fer+23], both assuming the extended Riemann hypothesis. These results assert that characteristic polynomials of random matrices behave, at least in terms of irreducibility, like random monic, integer polynomials, which have been long known to be irreducible with high probability [Wae34].

We also care about whether $\mathbb{Z}[\lambda] \simeq \mathbb{Z}[x]/(f(x))$ is the full ring of integers in $\mathbb{Q}(\lambda)$. In this regard, there are no results on characteristic polynomials of randomly selected integer matrices. There is only a conjecture of Wang and Yu asserting that symmetric integer matrices should make $\text{disc}f(x)$ squarefree (which implies $\mathbb{Z}[\lambda]$ is a ring of integers) a positive proportion of the time [WY16, Remark 1]. There is however, an analogous result for polynomials. Bhargava, Shankar, and Wang prove

that in any fixed degree, the proportion of monic, integer polynomials of bounded coefficient height with squarefree discriminant approaches

$$\frac{1}{2} \prod_{p \geq 3} \left(1 - \frac{3p-1}{p^2(p+1)}\right) \approx 0.307$$

as the height bound grows [BSW22]. They also prove that the proportion of polynomials that define a ring of integers approaches $\frac{6}{\pi^2} \approx 0.608$.

Like Wang and Yu, we conjecture that these positive-proportion results extend to the matrix setting. For $A \in \text{Mat}_n(\mathbb{Z})$ and $r > 0$, we use $B(A, r)$ denote the ℓ_∞ open ball of radius r in $\mathbb{R}^{n^2}$ around A .

Conjecture 4.1. *For $i = 1, 2$, there exists $p_i(n) > 0$ such that the proportion $P_i(A, r)$ of symmetric matrices in $B(A, r) \cap \text{Mat}_n(\mathbb{Z})$ with characteristic polynomial of squarefree discriminant (P_1) or trivial conductor (P_2) satisfies*

$$|P_i(A, r) - p_i(n)| < 2^{n^c} r^{-c'}$$

for any A and r and some absolute constants $c, c' > 0$. Both sequences $\{p_i(n)\}_n$ are bounded below by positive constants.

Factoring is the natural bottleneck for any test of Conjecture 4.1. Even if we considered matrices from an ℓ_∞ unit ball around the origin in $\mathbb{R}^{n^2}$, discriminants approach 900 bits when $n = 15$. For this reason, we present data only up to dimension 10 in Figure 2. In each dimension n , we test 5 centers A for $B(A, r)$ by choosing matrix entries uniformly from $(-2^{100/n^2}, 2^{100/n^2}) \cap \mathbb{Z}$ (insisting on matrix symmetry). The radius $r = 2^{50/n^2}$ defines a ball around each center from which we choose a random sample of 1000 symmetric integer matrices. The proportion of matrices whose characteristic polynomial has squarefree discriminant is on the left (one dot per sampled ball), and the proportion of those that define a full ring of integers is on the right.




FIGURE 2. Probability of squarefree discriminant (left) and trivial conductor.

Remark that when $n = 2$, it is straightforward to prove the exact limits of $P_1(A, r)$ and $P_2(A, r)$ as r grows. Both the discriminant formula of a symmetric matrix and the criterion for orders being maximal are far simpler when $n = 2$ compared to $n \geq 3$. Multiplying local probabilities gives

$$\lim_{r \rightarrow \infty} P_1(A, r) = \frac{1}{2} \prod_{p \equiv 3 \pmod{4}} \left(1 - \frac{1}{p^2}\right) \prod_{p \equiv 1 \pmod{4}} \left(1 - \frac{3p-2}{p^3}\right) \approx 0.366$$

and

$$\lim_{r \rightarrow \infty} P_2(A, r) = \frac{5}{4} \lim_{r \rightarrow \infty} P_1(A, r) \approx 0.458,$$

both uniformly in A .

For any n , the local factors appearing in the product should be bounded below (and above) by $1 - O(\frac{1}{p^2})$. Discriminant divisibility by p defines a codimension-1 subvariety in the space of symmetric matrices over $\mathbb{F}_p$. That is, roughly $\frac{1}{p}$ points in the subvariety meet this condition. For nonsingular points, lifting from mod p to mod p^2 imposes another $1/p$ density decrease, while the singular locus already has codimension 2 over $\mathbb{F}_p$. Hence the failure rate is $O(1/p^2)$. The author has not verified formally what local factors equal, but in the Euler product for $P_2(A, r)$ they appear to approach $1 - \frac{2}{p^2}$. Hence the predicted probability of

$$\prod_{p \text{ prime}} \left(1 - \frac{2}{p^2}\right) \approx 0.323$$

in Remark 1.2.

Finally, a note on the bound $2^{n^c} r^{-c'}$ in Conjecture 4.1: The denominator $r^{c'}$ matches the shape of Bhargava, Shankar, and Wang's error bound, in which $c' = \frac{1}{5} - \varepsilon$ [BSW22, Page 2 and Theorem 4.4]. (The value of c' balances errors from the sieve's tail and main body.) Regarding the numerator 2^{n^c} , such expressions hide behind big- O notation in [BSW22]. We care to be explicit, however, because convergence rate impacts algorithm running time. The shape of our bound is justified by identifying the fastest growing hidden constant in [BSW22], which appears to come from Bhargava's foundational work [Bha14]. In his proof of Theorem 3.3 (the quantitative Ekedahl sieve [Eke91]), the largest contributor to the final constant is a product of single-variable degrees of resultants built from an initial set of multivariate polynomials (see the third paragraph and the end of the proof of Lemma 3.1 in [Bha14]). In our setup (where Bhargava's " k " is 2 and the dimension is $\frac{1}{2}n(n+1)$), this product turns out to be $2^{O(n^3)}$.

Subroutine 2: Probabilistic method to adjust A so that $\det(xI_n - A)$ is irreducible and $\mathbb{Z}[\lambda]$ is the ring of integers in $\mathbb{Q}(\lambda)$.

Input: A symmetric $A \in \text{Mat}_n(\mathbb{Z})$ and a real number $\kappa > 1$

Output: A scaled, perturbed copy of A

```

1  $B \leftarrow A$  ▷ second copy of  $A$ 
2  $r \leftarrow (\sqrt{\kappa} - 1)/(n\|A^{-1}\|_2(\sqrt{\kappa} + 1))$ 
3  $\lambda \leftarrow$  any eigenvalue of  $A$ 
4 while  $[\mathbb{Q}(\lambda) : \mathbb{Q}] \neq n$  or  $1 \notin \mathfrak{f}(\mathbb{Z}[\lambda])$  do ▷ see Remark 4.2
5    $r \leftarrow 2r$ 
6    $A \leftarrow 2A$ 
7    $B \leftarrow$  symmetric with  $\max |(B - A)_{i,j}| < r$  ▷ chosen uniformly at random
8    $\lambda \leftarrow$  any eigenvalue of  $B$ 
9 return  $B, 1$ 
```

Remark 4.2. There are no known (non-quantum) polynomial time algorithms for testing whether $\mathbb{Z}[\lambda]$ is a ring of integers or has squarefree discriminant. By a result

of Lenstra [Len92, Theorem 4.4], these problems are polynomial-time equivalent to finding the largest square factor of an integer.

Theorem 4.3. *If Conjecture 4.1 is true, then Algorithm 1 with Subroutine 2 runs in bounded-error quantum polynomial time.*

Proof. The probability that B in line 7 satisfies the **while** loop condition is denoted $P_2(A, r)$ in Conjecture 4.1. The conjecture asserts that implies that the probability of success exceeds $\frac{1}{2}p_2(n)$ after **while** loop $\mathbb{Z}[\lambda]$ being the ring of integers in a degree- n number field is bounded below by $\frac{1}{2}$ after a constant number of iterations of Subroutine 2's **while** loop. This is exactly the criterion for terminating the **while** loop. To test whether $1 \in \mathfrak{f}(\mathbb{Z}[\lambda])$ in line 4, we use Lenstra's deterministic polynomial time (non-quantum) algorithm to reduce to the problem of computing the largest squarefree divisor of an integer [Len92, Theorem 4.4], which is solved in quantum polynomial time by Shor's factoring algorithm [Sho97].

All other steps in Algorithm 1 and Subroutine 2 evidently run in deterministic polynomial time in their input bit lengths, which are bounded by a polynomial in the bit length of M and κ . $\square$

Theorem 4.4. *From Algorithm 1 with Subroutine 2, $\alpha_1\beta, \dots, \alpha_n\beta$ is a $\mathbb{Z}$ -basis for an ideal $\mathfrak{b} \subseteq \mathbb{Z}[\lambda]$, the ring of integers in the totally real field $\mathbb{Q}(\lambda)$. Furthermore, if B is the matrix with i^{th} column $\Sigma(\alpha_i\beta)$ and A is its original value in line 2, then $AB^{-1}\Sigma(\mathfrak{b}) = M\mathbb{Z}^n$, and $\kappa_2(AB^{-1}) < \kappa$.*

Proof. There are two transformations in Algorithm 1 that we handle separately. The first is from (the lattices generated by) $A_1 := A$ to B_1 , where B_1 denotes the output of Subroutine 2. The second is from $A_2 := B_1$ to $B_2 := B$. We aim to bound each $\kappa_2(A_i B_i^{-1})$ by $\sqrt{\kappa}$, which will imply the claim:

$$\kappa_2(AB^{-1}) = \kappa_2((A_1 B_1^{-1})(A_2 B_2^{-1})) \leq \kappa_2(A_1 B_1^{-1}) \kappa_2(A_2 B_2^{-1}) < \kappa.$$

To bound $\kappa_2(A_1 B_1^{-1})$, let j denote the number of iterations that were needed in the **while** loop in Subroutine 2. If r is initial value in line 2, then

$$(4.1) \quad \|A_1 - 2^{-j} B_1\|_2 \leq n \max_{1 \leq i, j \leq n} |(A_1 - 2^{-j} B_1)_{i,j}| < nr = \frac{\sqrt{\kappa} - 1}{\|A_1^{-1}\|_2(\sqrt{\kappa} + 1)},$$

where the second inequality is line 7. By Lemma 3.6, we have

$$\sqrt{\kappa} > \kappa_2(A_1(2^{-j} B_1)^{-1}) = \kappa_2(A_1 B_1^{-1}).$$

Let E be the matrix whose i^{th} column is the i^{th} embedding of $(\alpha_1, \dots, \alpha_n)$, which are eigenvectors of A_2 , and let D be the corresponding diagonal matrix of eigenvalues. By comparing the columns corresponding to λ on either side of $A_2 E = ED$, we see that the α_i do indeed form a $\mathbb{Z}$ -basis for an ideal in $\mathbb{Z}[\lambda]$, which is the ring of integers in $\mathbb{Q}(\lambda)$ by line 4 of Subroutine 2.

Let C be the diagonal matrix with i^{th} diagonal entry a_i^{-1} . Since $A_2 = B_1$ is symmetric by line 7 of Subroutine 2, EC is an orthogonal matrix by the spectral theorem. Furthermore, $\mathbb{Q}(\lambda)$ is totally real.

Now let C' denote the diagonal matrix with the embeddings of β along its diagonal, and reassign r to be its value after line 8 of Algorithm 1. Observe that the output matrix is $B_2 = C' E^T$. We have

$$\kappa_2(A_2 B_2^{-1}) = \kappa_2(r(EC)^{-1} A_2 B_2^{-1}) \quad (\text{scaling/orthogonal transformation})$$

$$\begin{aligned}
&= \kappa_2(rCDE^T B_2^{-1}) && \text{by (3.1)} \\
&= \kappa_2(C(r^{-1}C')^{-1}) && \text{by definition of } B_2.
\end{aligned}$$

The distance between entries of $\mathbf{b}$ and $rV^{-1}\mathbf{a}$ is less than 1 by line 11. Hence the distance between entries of $\Sigma(\beta) = V\mathbf{b}$, which forms the diagonal of C' , and $V(rV^{-1}\mathbf{a})$, which forms the diagonal of rC , is less than $n\|V\| = n\lambda^{n-1}$ (because λ is the largest eigenvalue of A_2). Thus

$$(4.2) \quad \|C - r^{-1}C'\|_2 < r^{-1}n\lambda^{n-1} = \frac{\sqrt{\kappa} - 1}{\max_{i \leq n} |a_i|(\sqrt{\kappa} + 1)} = \frac{\sqrt{\kappa} - 1}{\|C^{-1}\|_2(\sqrt{\kappa} + 1)}.$$

By Lemma 3.6 again, $\kappa_2(C(r^{-1}C')^{-1}) < \sqrt{\kappa}$. $\square$

5. DETERMINISTIC SUBROUTINE TO FIND AN INVERTIBLE IDEAL

In this section, we provide a polynomial time subroutine that guarantees an irreducible characteristic polynomial for A from Algorithm 1 as well as an invertible output ideal. There is no need to consider general ℓ_p norms here, so p is often used to denote a prime number.

5.1. Polynomial irreducibility. The simplest way to obtain an irreducible characteristic polynomial is to fix a target irreducible polynomial $f(x)$ modulo some small prime p and find a symmetric matrix S_p with $\text{char}_{S_p}(x) \equiv f(x) \pmod{p}$. We can enforce that the final output matrix A produced by the subroutine satisfies $A \equiv S_p \pmod{p}$. We briefly review the known efficient techniques to achieve this with $p = 2$.

Shoup's deterministic algorithm finds an irreducible $f(x) \in \mathbb{F}_2[x]$ of degree n [Sho90]. (The symmetry constraint makes it hard to use polynomials that are Eisenstein at p .) Let C_f be its companion matrix satisfying $f(x) = x^n + \sum_i (C_f)_{i+1,n} x^i$.

A standard technique (see [BPR06], for example) to produce a symmetric (Hankel) matrix H satisfying $HC_f = C_f^T H$ is letting $H_{i,j} = \text{Tr}(x^{i+j-2})$, the trace from $\mathbb{F}_2[x]/(f(x))$ to $\mathbb{F}_2$. In other words, H represents the trace pairing with respect to the basis $1, x, \dots, x^{n-1}$. For example, when $f(x) = x^3 + x + 1$, the traces of $1, x, x^2, x^3$, and x^4 are $1, 0, 0, 1$, and 0 (the antidiagonals of H below). And indeed,

$$HC_f = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} = C_f^T H.$$

This identity holds because C_f represents multiplication by x in $\mathbb{F}_2[x]/(f(x))$, which is self-adjoint with respect to the trace bilinear form. The equation above simply expresses that the trace is self-adjoint with respect to multiplication by x (as detailed in the proof of Lemma 5.1).

The trace pairing is a non-alternating form. By Dickson's classification of symmetric bilinear forms (including over $\mathbb{F}_2$) [Dic01], this implies the trace pairing is diagonalizable, meaning $P^T H P = D$ for some diagonal matrix D and change of basis matrix P . Over $\mathbb{F}_2$, $D = I_n$ is forced, so $H = (P P^T)^{-1}$. Thus $HC_f = C_f^T H$ reorganizes to $P^{-1} C_f P = (P^{-1} C_f P)^T$. We have found a symmetric matrix with characteristic polynomial $f(x)$. In the example, $1 + x, 1 + x^2, 1 + x + x^2$ is an

orthonormal basis with respect to the trace form, and indeed,

$$P^{-1}C_fP = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

is a symmetric matrix with characteristic polynomial $x^3 + x + 1$.

Over fields of characteristic other than 2, the Gram–Schmidt algorithm diagonalizes H . Over fields of characteristic 2, non-degenerate does not imply non-alternating, so Gram–Schmidt alone can fail if it stumbles into a subspace in which the bilinear form is alternating. Still, there are simple algorithms to bypass this obstacle. See, for example, the StackExchange post by GreginGre [mat20], which is summarized in the last paragraph of the proof below.

Lemma 5.1. *A symmetric matrix in $S_2 \in \text{GL}_n(\mathbb{F}_2)$ with irreducible characteristic polynomial can be found in deterministic polynomial time.*

Proof. Shoup’s algorithm finds an irreducible $f(x) \bmod 2$ of degree n in polynomial time. The matrix symmetric H representing the trace pairing from $\mathbb{F}_2[x]/(f(x))$ to $\mathbb{F}_2$ with respect to the basis $1, x, \dots, x^{n-1}$ can be computed in $O(n^2)$ bit operations. Indeed, if c_i is the coefficient of x^{n-i} in $f(x)$, Newton’s formulas assert

$$\text{Tr}(x^k) = kc_k + \sum_{i=1}^{k-1} c_{k-i} \text{Tr}(x^i).$$

So the $2n - 1$ antidiagonals of H can be computed recursively, each in $O(n)$ bit operations.

Observe that

$$\begin{aligned} (HC_f)_{i,j} &= \sum_{k=1}^n \text{Tr}(x^{i+k-2})(C_f)_{k,j} && \text{by definition of } H \\ &= \text{Tr}\left(x^{i-1} \sum_{k=1}^n (C_f)_{k,j} x^{k-1}\right) && \text{by bilinearity of the trace} \\ &= \text{Tr}(x^{i-1} \cdot (x \cdot x^{j-1})) && \text{since } C_f \text{ represents multiplication by } x \\ &= \text{Tr}(x^{j-1} \cdot (x \cdot x^{i-1})) \\ &= \text{Tr}\left(x^{j-1} \sum_{k=1}^n (C_f)_{k,i} x^{k-1}\right) && \text{since } C_f \text{ represents multiplication by } x \\ &= \sum_{k=1}^n (C_f^T)_{i,k} \text{Tr}(x^{j+k-2}) && \text{by bilinearity of the trace} \\ &= (C_f^T H)_{i,j} && \text{by definition of } H. \end{aligned}$$

In particular, diagonalizing H with some matrix P completes the process as previously described—if $P^T H P = I_n$, then $S_2 := P^{-1}C_fP$ is symmetric.

Since finite field extensions are separable, the trace pairing is a non-degenerate. Furthermore, the traces of $1 \cdot 1, x \cdot x, \dots, x^{n-1} \cdot x^{n-1}$ cannot all vanish because these monomials are a basis for $\mathbb{F}_2[x]/(f(x))$ (apply the Frobenius automorphism to $1, x, \dots, x^{n-1}$). In other words, the trace pairing is nonalternating.

To find a diagonalizing basis $g_1(x), \dots, g_n(x)$, we follow [mat20]: Let $V_n = \mathbb{F}_2[x]/(f(x))$, and pick any $g(x) \in V$ such that $\text{Tr}(g(x)^2) = 1$. If the trace form remains nonaltering on $\langle g(x) \rangle^\perp$, let $g_1(x) := g(x)$ and $V_{n-1} := \langle g_1(x) \rangle^\perp$. Otherwise, pick any $h_1(x), h_2(x) \in \langle g(x) \rangle^\perp$ such that $\text{Tr}(h_1(x)h_2(x)) = 1$ (such polynomials exist since the trace is nondegenerate on $\langle g(x) \rangle^\perp$), and set $g_1(x) := g(x) + h_1(x) + h_2(x)$ and $V_{n-1} := \langle g_1(x) \rangle^\perp$. Note that $g(x) + h_1(x) \in \langle g_1(x) \rangle^\perp$ and $\text{Tr}((g(x) + h_1(x))^2) = 1$, allowing the procedure to continue on V_{n-1} . $\square$

5.2. Ideal invertibility. This is the primary challenge of the deterministic subroutine for Algorithm 1. The ring at hand is $\mathcal{O} := \mathbb{Z}[x]/(f(x))$, where $f(x) = \det(xI_n - A)$, and the ideal that must be invertible is generated by the last column entries of $\text{adj}(xI_n - A)$. Our strategy is to make this ideal coprime to $f'(x)$, which is sufficient but not necessary for invertibility (see Theorem 5.12). This turns out to be possible through a symmetric perturbation of A that is small enough to negligibly affect the shape of the lattice. Specifically, our perturbation takes the form

$$(5.1) \quad A + \begin{bmatrix} & a_2 & & \\ a_2 & b_2 & \ddots & \\ & \ddots & \ddots & a_n \\ & & a_n & b_n \end{bmatrix}.$$

Definition 5.2. Let A_j denote the top-left $j \times j$ minor of A , let $f_j(x) = \det(xI_j - A_j)$, let $\mathfrak{a}_j \subseteq \mathbb{Z}[x]$ denote the ideal generated by the last column entries of $\text{adj}(xI_j - A_j)$, and call the top entry in this column $g_{j-1}(x)$.

The search for a_j and b_j that make $(\mathfrak{a}_n, f'_n(x)) = \mathbb{Z}[x]$ is performed sequentially. Assume $a_2, b_2, \dots, a_{j-1}, b_{j-1}$ have already been chosen so that $(\mathfrak{a}_{j-1}, f'_{j-1}(x), p) = \mathbb{Z}[x]$ for any prime p . We first choose a_j to ensure that $(\mathfrak{a}_j, f'_{j-1}(x), p) = \mathbb{Z}[x]$ for all p . To see why this is possible, observe that the top $j-1$ generators of $\mathfrak{a}_j$ can be computed as the product

$$(5.2) \quad \text{adj}(xI_{j-1} - A_{j-1}) \begin{bmatrix} A_{1,j} \\ \vdots \\ A_{j-1,j} + a_j \end{bmatrix}.$$

For any prime $\mathfrak{p} \mid p$ in $\mathbb{Z}[x]$ that contains $f'_{j-1}(x)$, we have assumed that the final entry in at least one row of $\text{adj}(xI_{j-1} - A_{j-1})$ (these final entries are the coefficients of a_j above) is not in $\mathfrak{p}$. Therefore, there is at most one bad congruence class $a_j \bmod p$ associated to $\mathfrak{p}$, where $p \in \mathbb{Z}$ is the prime below $\mathfrak{p}$. We simply increment a_j until it avoids every bad congruence class from the finitely many ramified primes with odd norm in $\mathbb{Z}[x]/(f_{j-1}(x))$.

Once a_j is chosen, the purpose of b_j is to switch rings; to complete the induction, we need $\mathfrak{a}_j$ to avoid ramified primes with odd norm in $\mathbb{Z}[x]/(f_j(x))$, not in $\mathbb{Z}[x]/(f_{j-1}(x))$. By cofactor expansion along the bottom row of $xI_j - A_j$,

$$f_j(x) = (x - A_{j,j} - b_j)f_{j-1}(x) + \sum_{i=1}^{j-1} A_{j,i} \text{adj}(xI_j - A_j)_{i,j}.$$

This is a linear combination of the generators of $\mathfrak{a}_j$. Hence, if $\mathfrak{p}$ is a prime in $\mathbb{Z}[x]$ that contains $\mathfrak{a}_j$, then $\mathfrak{p}$ automatically contains $f_j(x)$. The job of b_j is to make

sure $\mathfrak{p}$ does not also contain the derivative $f'_j(x)$, lest it be ramified in $\mathbb{Z}[x]/(f_j(x))$. Since

$$f'_j(x) = (x - A_{j,j} - b_j)f'_{j-1}(x) - (A_{j,j} + b_j)f_{j-1}(x) + \sum_{i=1}^{j-1} A_{j,i} \text{adj}(xI_j - A_j)'_{i,j},$$

and since a_j was specifically chosen so that $f'_{j-1}(x)$ is not in $\mathfrak{p}$, there is at most one bad congruence class $b_j \bmod p$ associated to $\mathfrak{p}$. As before, we increment b_j until it avoids every bad congruence class from the finitely many odd-norm primes dividing $\mathfrak{a}_j$. This completes the inductive step.

There are two potential issues that can arise in the search for a_j or b_j . First, since $f_{j-1}(x)$ can have up to $j-1$ roots $\bmod p$, if $p < j$, it is possible that there is no acceptable congruence class $\bmod p$ for a_j or b_j ; the search would never terminate. In particular, we must control A modulo every prime less than n . In fact, we control A at some larger primes as well.

Definition 5.3. For $p > 2$, define $S_p \in \text{Mat}_n(\mathbb{F}_p)$ to be

$$\begin{bmatrix} 1 & 1 & & \\ 1 & & \ddots & \\ & \ddots & & 1 \\ & & 1 & \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} & 1 & & \\ 1 & & \ddots & \\ & \ddots & & 1 \\ & & 1 & \end{bmatrix},$$

the former if $p \mid n+1$, the latter if $p \nmid n+1$. Let q be the product of all primes less than $4n^2$, and fix some $S_q \in \text{Mat}_n(\mathbb{Z})$ with maximum entry magnitude at most $\frac{q}{2}$ and $S_q \equiv S_p \bmod p$ for all $p \mid q$ (even $p = 2$ using S_2 from Lemma 5.1).

Lemma 5.4. *The discriminant of characteristic polynomial of S_q is coprime to q .*

Proof. The discriminant of $\det(xI_n - S_q)$ reduced $\bmod p$ equals that of $\det(xI_n - S_p)$. When $p = 2$, this polynomial is irreducible by Lemma 5.1, and it therefore has a nonzero discriminant in $\mathbb{F}_2$ (a perfect field).

For odd p that do not divide $n+1$, $\det(xI_n - S_p)$ is the Chebyshev polynomial of the second kind, denoted $U_n(x/2)$, which are defined by the $U_0(x/2) = 1$, $U_1(x/2) = x$, and $U_n(x/2) = xU_{n-1}(x/2) - U_{n-2}(x/2)$. The well-known substitution $x = y + y^{-1}$ makes

$$U_n(x/2) = \frac{y^{n+1} - y^{-n-1}}{y - y^{-1}} = \frac{y^{-n}(y^{2n+2} - 1)}{y^2 - 1}.$$

In particular, the roots of $U_n(x/2)$ are $x = \zeta^i + \zeta^{-i}$ for $i = 1, \dots, n$ and ζ a primitive $(2n+2)^{\text{th}}$ root of unity. These roots are distinct in $\overline{\mathbb{F}}_p$ when $p \nmid n+1$.

For odd p that divide $n+1$, $\det(xI_n - S_p) = U_n(x/2) - U_{n-1}(x/2)$. Using $x = y + y^{-1}$ again, this becomes

$$\frac{(y^{n+1} - y^{-n-1}) - (y^n - y^{-n})}{y - y^{-1}} = \frac{y^{-n}(y^{2n+1} + 1)}{y + 1}.$$

Thus the roots are $x = \zeta^{2i-1} + \zeta^{-2i+1}$ for $i = 1, \dots, n$ and ζ a primitive $(4n+2)^{\text{th}}$ root of unity. These roots are distinct in $\overline{\mathbb{F}}_p$ when $p \nmid 2n+1$, which is true when $p \mid n+1$. $\square$

Remark 5.5. It follows from this proof that odd primes $p < 4n^2$ split completely in the ring $\mathbb{Z}[\lambda]$ produced in Algorithm 1. Alternative choices for S_p in Definition 5.3

can produce whatever splitting behavior we prefer. (Insisting that S_p is symmetric is not strictly necessary. It just makes our analysis of Algorithm 1 cleaner.)

The reason for the additional restriction at primes between n and $4n^2$ stems from the second issue: the search interval required to find a_j and b_j could be too large. While the total number of congruence classes modulo primes that a_j and b_j must avoid is easily bounded by a polynomial in $n \log \|M\|$, sieve theory requires a stronger hypothesis. The sifting density, call it ρ (unfortunately “ κ ” in the literature, which we reserve for the condition number), is an upper bound on the average number of congruence classes $w(p)$ to be avoided per prime p in any interval:

$$\sum_{y < p < z} \frac{w(p)}{p} = \rho \log \left(\frac{\log z}{\log y} \right) + O(1)$$

for any $z > y \geq 2$. Enforcing $w(p) = 0$ when $p < 4n^2$ significantly reduces ρ in our setup. Still, as soon as the interval bound y reaches $4n^2$, $w(p)$ can spike up to j if we are unlucky. Since sieves place ρ in the exponent of our search intervals for a_j and b_j , this is a big problem for the desired polynomial running time. We might overcome this exponent by distributing the huge search interval over all j entries in the final column of A_j rather than just the bottom two entries, but searching through a j -dimensional box is also inherently exponential in running time.

The solution is to impose an additional constraint on a_j , asking that it “prepares” a nice sieve problem for b_j . Similarly, b_j must prepare the sieve for a_{j+1} . Recall that b_j avoids congruence classes at primes dividing $[\mathbb{Z}[x] : \mathfrak{a}_j]$. If $[\mathbb{Z}[x] : \mathfrak{a}_j]$ were a randomly selected integer from some large interval, the expected sifting density would be $\kappa = 0$, meaning

$$(5.3) \quad \sum_{p \text{ prime}} \frac{1}{p} v_p([\mathbb{Z}[x] : \mathfrak{a}]) = O(1),$$

where v_p is the valuation at p . This is precisely the additional constraint on a_j . Similarly, since a_{j+1} must avoid congruence classes at primes dividing $\text{disc}f(x)$, we ask b_j to enforce

$$(5.4) \quad \sum_{p \text{ prime}} \frac{1}{p} v_p(\text{disc}f(x)) = O(1).$$

The cost is a small addition to one search interval in exchange for a good bound on the next search interval.

As a final remark, the sums in (5.3) and (5.4) cannot be computed exactly because there is no known polynomial time algorithm for factoring integers. In line of Subroutine 2, we define a bound r beyond which the primes contribute negligibly.

5.3. The algorithm. We use the following standard notation:

Definition 5.6. The *radical* of an ideal $\mathfrak{a}$ in a commutative ring $\mathcal{O}$ is

$$\sqrt{\mathfrak{a}} := \{\alpha \in \mathcal{O} : \alpha^k \in \mathfrak{a} \text{ for some } k \geq 0\}.$$

An alternative phrasing of $2 \in \sqrt{\mathfrak{a}}$ for some finite index ideal $\mathfrak{a} \subseteq \mathcal{O}$ (as in lines 6 and 10 below, where $\mathcal{O} = \mathbb{Z}[x]$) is $[\mathcal{O} : \mathfrak{a}] = 2^k$ for some $k \geq 0$.

Definition 5.7. Let $\mathcal{O}$ be a commutative ring. The *valuation* of an ideal $\mathfrak{a} \subseteq \mathcal{O}$ at the prime $\mathfrak{p} \subseteq \mathcal{O}$, denoted $v_{\mathfrak{p}}(\mathfrak{a})$, is the largest integer k such that $\mathfrak{a} \subseteq \mathfrak{p}^k$.

Subroutine 3: Deterministic method to adjust A so that $\det(xI_n - A)$ is irreducible and the output ideal of Algorithm 1 is invertible.

Input: A symmetric $A \in \text{Mat}_n(\mathbb{Z})$ and a real number $\delta > 1$
Output: A scaled, perturbed copy of A

```

1  $q' \leftarrow n^2 q \lceil \|A^{-1}\|_2(\sqrt{\kappa} + 1)/(\sqrt{\kappa} - 1) \rceil$  ▷ see Definition 5.3
2  $A \leftarrow 3q' \lceil \log(q' \|A\|_2) \rceil A + S_q$ 
3  $r \leftarrow n^2(\log \|A\|_2)$ 
4 for  $j$  from 2 to  $n$  do
5    $s \leftarrow \sum_p \frac{1}{p} v_p([\mathbb{Z}[x] : \mathfrak{a}_j])$  for  $4n^2 < p < r$  ▷ see Definitions 5.2 and 5.7
6   while  $2 \notin \sqrt{(\mathfrak{a}_j, f'_{j-1}(x))}$  or  $s > \frac{j}{3n}$  do ▷ see Definition 5.6
7      $A_{j,j-1}, A_{j-1,j} \leftarrow A_{j,j-1} + q$ 
8     recompute  $\mathfrak{a}_j$  and  $s$ 
9      $s \leftarrow \sum_p \frac{1}{p} v_p(\text{disc } f_j(x))$  for  $4n^2 < p < r$ 
10    while  $2 \notin \sqrt{(\mathfrak{a}_j, f'_j(x))}$  or  $s > \frac{j}{3n}$  do
11       $A_{j,j} \leftarrow A_{j,j} + q$ 
12      recompute  $f_j(x)$  and  $s$ 
13 return  $A, r$ 

```

The following lemma is known, in that it is an immediate consequence of established results. A proof is included for completeness.

Lemma 5.8. *For all $x \geq 2$, we have for primes p that*

$$\sum_{p>x} \frac{1}{p^2} < \frac{1}{x \log x} - \frac{1}{3x \log^2 x}.$$

Proof. We evaluate the Riemann–Stieltjes integral with integration by parts:

$$\sum_{p>x} \frac{1}{p^2} = \int_x^\infty \frac{1}{t^2} d\pi(t) = -\frac{\pi(x)}{x^2} + \int_x^\infty \frac{2\pi(t)}{t^3} dt.$$

(A weaker version of) Dusart’s inequality for $x \geq 599$, [Dus18, Corollary 5.2],

$$\frac{x}{\log x} \left(1 + \frac{1}{\log x}\right) < \pi(x) < \frac{x}{\log x} \left(1 + \frac{4}{3 \log x}\right),$$

then bounds our initial sum from above. We apply integration by parts twice to the result:

$$\begin{aligned} & -\frac{1}{x \log x} - \frac{1}{x \log^2 x} + \int_x^\infty \left(\frac{2}{t^2 \log t} + \frac{8}{3t^2 \log^2 t} \right) dt \\ & = \frac{1}{x \log x} - \frac{1}{3x \log^2 x} - \int_x^\infty \frac{4}{3t^2 \log^3 t} dt. \end{aligned}$$

The final integral is positive, so the proof is complete for $x \geq 599$. The bound is easily checked for $2 \leq x < 599$. $\square$

Lemma 5.9. *Let r be as in line 3 and recall notation from (5.1). The perturbations $\mathfrak{a}_j$ and b_j found in the **while** loops of Subroutine 3 are bounded by qr .*

Proof. Let t denote the value of $\|A\|_2$ used in line 3. From Definition 5.3, $q = e^{\vartheta(4n^2)}$, where ϑ is Chebyshev’s theta function. (A weaker version of) Rosser and

Schoenfield's lower bound [RS62, Corollary (3.16)] gives $e^{\vartheta(4n^2)} > e^{3n^2}$ whenever $n \geq 4$. Combining this with lines 1 and 2 crudely gives

$$(5.5) \quad t \geq 3q' \log(q') - \|S_q\|_2 \geq 3n^2 q \log(n^2 q) - \frac{nq}{2} > 9n^4 e^{3n^2}.$$

whenever $n \geq 4$. That this lower bound on t also holds when $n = 2$ or 3 is easily checked. The penultimate expression above also shows (crudely again) that

$$(5.6) \quad q < \frac{t}{2n^2 \log t}.$$

To simplify the exposition, we first bound $|\text{disc } f_j(x)|$ assuming $0 \leq a_i, b_i \leq qr$ for all $i \leq j$. The spectral norm of any minor is bounded above by that of the matrix, and there are at most two values of a_i and b_i in any row or column (see (5.1)), so all minors have spectral norm at most $t + 2qr$. A minor's determinant of A_j is bounded by its spectral norms raised to the power of its dimensions. Thus the coefficient of x^{j-i} in $f_j(x)$ is bounded by $\binom{j}{i}(t + 2qr)^i$, and the coefficient of x^{j-1-i} in $f'_j(x)$ is bounded by $\binom{j}{i}(j-i)(t + 2qr)^i < \binom{j}{i}2^{-3j}(t + 2qr)^{i+1}$. Hence the sum of all coefficient magnitudes of $f_j((t + 2qr)x)$ is bounded by

$$\sum_{i=0}^j \binom{j}{i} (t + 2qr)^i (t + 2qr)^{j-i} = 2^j (t + 2qr)^j.$$

Similarly, the sum of coefficient magnitudes of $f'_j((t + 2qr)x)$ is bounded by $2^{-2j}(t + 2qr)^j$. These sums bound the ℓ_2 lengths of the columns of the Sylvester matrix for $f_j((t + 2qr)x)$ and $f'_j((t + 2qr)x)$, whose determinant is bounded by Hadamard's inequality in the third line below:

$$\begin{aligned} |\text{disc } f_j(x)| &= |\text{res}(f_j(x), f'_j(x))| \\ &= \frac{1}{(t + 2qr)^{j(j-1)}} |\text{res}(f_j((t + 2qr)x), f'_j((t + 2qr)x))| \\ &\leq \frac{2^{j(j-1)-2j(j)} (t + 2qr)^{j(j+1-1)}}{(t + 2qr)^{j(j-1)}} \\ &< (t + qr)^{j^2}. \end{aligned}$$

Now, from (5.6) and the definition of r , $t + 2qr$ is bounded by $t(1 + 1/\log n)$. In particular, if $j < n$ we have

$$(5.7) \quad \frac{\log |\text{disc } f_j(x)|}{r \log n} < \frac{(n-1)^2 \log(t(1 + 1/\log n))}{n^2 \log t \log n}.$$

This is a decreasing function of t . Substituting the lower bound on t from (5.5) into (5.7) gives a function of n that is less than $\frac{4}{9}$ for all $n \geq 2$. Therefore,

$$(5.8) \quad \log |\text{disc } f_j(x)| < \frac{4}{9} r \log n.$$

Now fix some **for** loop iteration $2 \leq j \leq n$, and assume for induction that $a_i, b_i \leq qr$ for all $i < j$. Consider the **while** loop in line 6. Forcing $2 \in \sqrt{(\mathbf{a}_j, f'_{j-1}(x))}$ is equivalent to $(\mathbf{a}_j, f'_{j-1}(x), p) = \mathbb{Z}[x]$ for every odd prime p . If $p < 4n^2$, then $\mathbf{a}_j$ contains the element $g_{j-1}(x)$ defined in Definition 5.2. Since $A \equiv S_p \pmod p$ by line 2, we see from Definition 5.3 that the bottom-left $(j-1) \times (j-1)$ minor determinant of $xI_j - A_j$, called $g_{j-1}(x)$ as per Definition 5.2, is equivalent to $1 \pmod p$. Hence $(\mathbf{a}_j, f'_{j-1}(x), p) = \mathbb{Z}[x]$ is automatic when $p < 4n^2$.

Suppose $p > 4n^2$. If $j \geq 3$, line 10 from the previous **for** loop iteration guarantees

$$(5.9) \quad (\mathfrak{a}_{j-1}, f'_{j-1}(x), p) = \mathbb{Z}[x].$$

We also have

$$\begin{aligned}
 \sum_{p > 4n^2} \frac{1}{p} v_p(\text{disc } f_{j-1}(x)) &= \sum_{4n^2 < p < r} \frac{1}{p} v_p(\text{disc } f_{j-1}(x)) + \sum_{p \geq r} \frac{1}{p} v_p(\text{disc } f_{j-1}(x)) \\
 &\leq \frac{j-1}{3n} + \sum_{p \geq r} \frac{1}{p} v_p(\text{disc } f_{j-1}(x)) \quad \text{by line 10 for } j-1 \\
 &< \frac{1}{3} + \frac{1}{r \log r} \sum_{p \geq r} v_p(\text{disc } f_{j-1}(x)) \log p \\
 &\leq \frac{1}{3} + \frac{\log |\text{disc } f_{j-1}(x)|}{r \log r} \\
 &< \frac{1}{3} + \frac{\log |\text{disc } f_{j-1}(x)|}{r \log(n^2 \log e^{3n^2})} \quad \text{by (5.5)} \\
 &< \frac{1}{3} + \frac{\log |\text{disc } f_{j-1}(x)|}{4r \log n} \\
 (5.10) \quad &< \frac{4}{9} \quad \text{by (5.8)}.
 \end{aligned}$$

Properties (5.9) and (5.10) also hold when $j = 2$. Indeed, $f_1(x) = x - A_{1,1}$, so $(\mathfrak{a}_1, f'_1(x), p) = (\mathfrak{a}_1, 1, p) = \mathbb{Z}[x]$ and $\sum_p \frac{1}{p} v_p(\text{disc } f(x)) = \sum_p \frac{1}{p} v_p(1) = 0$.

Let $\mathfrak{p}$ be the pullback to $\mathbb{Z}[x]$ of a ramified in $\mathbb{Z}[x]/(f_{j-1}(x))$ above p , meaning $\mathfrak{p}$ is generated by p and a common irreducible factor of $f_{j-1}(x) \bmod p$ and $f'_{j-1}(x) \bmod p$. As we checked in (5.9), such a prime does not contain $\mathfrak{a}_{j-1}$. By computing the first $j-1$ generators of $\mathfrak{a}_j$ as in (5.2), we see that $\mathfrak{a}_{j-1} \not\subseteq \mathfrak{p}$ implies there is at most one congruence class $a_j \bmod p$ that makes $\mathfrak{a}_j \subseteq \mathfrak{p}$. Since $v_p(\text{disc } f_{j-1}(x))$ is an upper bound on the number of ramified primes above p , we conclude that the total number of congruence classes $\bmod p$ that a_j must avoid, call this number $w(p)$, is at most $v_p(\text{disc } f_{j-1}(x))$.

As a_j increments through the sequence $0, q, \dots, [r]q$, it lands on one of the $w(p)$ congruence classes $\bmod p$ at most $w(p) \left(\frac{r}{p} + 1\right)$ times. So altogether, the total number of values of a_j in this sequence of length $[r] + 1$ for which $(\mathfrak{a}_j, f'_{j-1}(x))$ is contained in a prime above some $p > 4n^2$ is at most

$$\begin{aligned}
 \sum_{p > 4n^2} w(p) \left(\frac{r}{p} + 1\right) &\leq r \sum_{p > 4n^2} \frac{v_p(\text{disc } f_{j-1}(x))}{p} + \sum_{p > 4n^2} v_p(\text{disc } f_{j-1}(x)) \\
 &\leq \frac{4r}{9} + \sum_{p > 4n^2} v_p(\text{disc } f_{j-1}(x)) \quad \text{by (5.10)} \\
 &< \frac{4r}{9} + \frac{1}{\log 4n^2} \sum_{p > 4n^2} v_p(\text{disc } f_{j-1}(x)) \log p \\
 &< \frac{4r}{9} + \frac{\log |\text{disc } f_{j-1}(x)|}{2 \log n} \\
 (5.11) \quad &< \frac{2r}{3} \quad \text{by (5.7)}.
 \end{aligned}$$

Next we count how many choices of a_j among $0, q, \dots, \lfloor r \rfloor q$ fail line 6 because they make $\sum_p \frac{1}{p} v_p([\mathbb{Z}[x] : \mathfrak{a}_j])$ too large for primes $4n^2 < p < r$. Since $\mathfrak{a}_j$ contains $f_{j-1}(x)$ for any choice of a_j , we only ever consider primes $\mathfrak{p} \subseteq \mathbb{Z}[x]$ generated by $p > 4n^2$ together with some irreducible factor of $f_{j-1}(x) \bmod p$. Given $\mathfrak{p}$, let $w_{\mathfrak{p}}$ denote the product of the inertial degree, call it $d(\mathfrak{p}|p)$ (the degree of the corresponding irreducible factor of $f_{j-1}(x) \bmod p$), with the minimal valuation of $\mathfrak{a}_j$ at $\mathfrak{p}$ as a_j runs through the sequence $0, q, \dots, \lfloor r \rfloor q$. Also let $w'_{\mathfrak{p}}(a)$ denote the value of $d(\mathfrak{p}|p)v_{\mathfrak{p}}(\mathfrak{a}_j) - w_{\mathfrak{p}}$ when $\mathfrak{a}_j$ is determined by $a_j = aq$. This way,

$$(5.12) \quad \sum_{\mathfrak{p}|p} (w_{\mathfrak{p}} + w'_{\mathfrak{p}}(a)) = v_p([\mathbb{Z}[x] : \mathfrak{a}_j]).$$

The coefficients of a_j in (5.2) all lie in $\mathfrak{a}_{j-1}$, and thus they lie in $\mathfrak{p}^k$ with $k = v_{\mathfrak{p}}(\mathfrak{a}_{j-1})$. In particular, changing a has no effect on the generators of $\mathfrak{a}_j$ reduced mod $\mathfrak{p}^k$. However, since $\mathfrak{a}_{j-1} \not\subseteq \mathfrak{p}^{k+1}$, there is at most one choice of $a_j \bmod p$ for which $v_{\mathfrak{p}}(\mathfrak{a}_j)$ could possibly exceed k . In other words, $w_{\mathfrak{p}} \leq d(\mathfrak{p}|p)k$. Thus for any choice of a_j ,

$$(5.13) \quad \begin{aligned} \sum_{4n^2 < p < r} \sum_{\mathfrak{p}|p} \frac{w_{\mathfrak{p}}}{p} &\leq \sum_{4n^2 < p < r} \sum_{\mathfrak{p}|p} \frac{1}{p} d(\mathfrak{p}|p) v_{\mathfrak{p}}(\mathfrak{a}_{j-1}) \\ &= \sum_{4n^2 < p < r} \frac{1}{p} v_p([\mathbb{Z}[x] : \mathfrak{a}_{j-1}]) \\ &\leq \frac{j-1}{3n} \quad \text{by line 6 for } j-1. \end{aligned}$$

Furthermore, we know $w'_{\mathfrak{p}}(a)$ is positive at most once for $a = 0, \dots, p-1$, and its value cannot exceed $d(\mathfrak{p}|p)v_{\mathfrak{p}}(f_{j-1}(x))$ because $f_{j-1}(x) \in \mathfrak{a}_j$. But this product is simply the degree of the largest power of the irreducible factor defining $\mathfrak{p}$ that divides $f_{j-1}(x) \bmod p$. Since $f_{j-1}(x)$ is monic of degree $j-1$, we have

$$(5.14) \quad \sum_{\mathfrak{p}|p} \sum_{a=0}^{p-1} w'_{\mathfrak{p}}(a) \leq j-1.$$

The number of times a hits any given congruence class mod p in the sequence $0, 1, \dots, \lfloor r \rfloor$ is at most $\frac{r}{p} + 1$, allowing us to bound the average of the double sum denoted $S(a)$ below:

$$(5.15) \quad \begin{aligned} \frac{1}{\lfloor r \rfloor + 1} \sum_{a=0}^{\lfloor r \rfloor} \underbrace{\sum_{4n^2 < p < r} \sum_{\mathfrak{p}|p} \frac{w'_{\mathfrak{p}}(a)}{p}}_{S(a)} &< \sum_{4n^2 < p < r} \frac{1}{rp} \left(\frac{r}{p} + 1 \right) \sum_{\mathfrak{p}|p} \sum_{a=0}^{p-1} w'_{\mathfrak{p}}(a) \\ &\leq (j-1) \sum_{4n^2 < p < r} \left(\frac{1}{p^2} + \frac{1}{rp} \right) \quad \text{by (5.14)} \\ &< (j-1) \left(\sum_{p > 4n^2} \frac{1}{p^2} + \frac{1}{r} \sum_{4n^2 < p < r} \frac{1}{p} \right). \end{aligned}$$

Setting $x = 4n^2$ in the statement of Lemma 5.8 bounds the first summation above, and the second is bounded according to Rosser and Schoenfeld by $\log \log r - 1 + (\log r)^{-2}$ [RS62, Corollary (3.20)]. (Note that the constant -1 in our bound is less than in [RS62, Corollary (3.20)]; this is more than accounted for by our lack of

primes less than $4n^2$.) Using $r = n^2 \log t \geq 3n^4$ by (5.5), we see that the average value of $S(a)$ is at most

$$(5.16) \quad (n-1) \left(\frac{1}{4n^2 \log 4n^2} - \frac{1}{12n^2 \log^2 4n^2} + \frac{\log \log 3n^4 - 1 + (\log 3n^4)^{-2}}{3n^4} \right).$$

This is less than $\frac{1}{9n}$ for all $n \geq 2$.

No more than one-third of all a can make $S(a)$ at least thrice the average. For the remaining choices of a , we have $S(a) < \frac{1}{3n}$, giving

$$\begin{aligned} \sum_{4n^2 < p < r} \frac{1}{p} v_p([\mathbb{Z}[x] : \mathbf{a}_j]) &= \sum_{4n^2 < p < r} \sum_{\mathfrak{p} | p} \frac{1}{p} (w_{\mathfrak{p}} + w'_{\mathfrak{p}}(a)) && \text{by (5.12)} \\ &\leq \frac{j-1}{3n} + S(a) && \text{by (5.13)} \\ &< \frac{j}{3n}, \end{aligned}$$

just as line 6 demands.

Combining this with (5.11) shows that the total number of failures among $a_j = 0, q, \dots, \lfloor r \rfloor q$ for either condition in line 6 is less than $\frac{r}{3} + \frac{2r}{3} < \lfloor r \rfloor + 1$. In particular, at least one success is guaranteed, so $a_j < qr$.

The proof that $b_j < qr$ is nearly identical, but with $[\mathbb{Z}[x] : \mathbf{a}_j]$ now playing the role of $\text{disc} f_{j-1}(x)$, and $\text{disc} f_j(x)$ now playing the role of $[\mathbb{Z}[x] : \mathbf{a}_j]$. Let us highlight the minor adjustments.

To start, (5.8) can be deduced for $[\mathbb{Z}[x] : \mathbf{a}_{j+1}]$ in place of $\text{disc} f_j(x)$ by bounding the resultant of $f_j(x)$ and $g_j(x)$ rather than $f_j(x)$ and $f'_j(x)$. This is because

$$|\text{res}(f_j(x), g_j(x))| = [\mathbb{Z}[x] : (f_j(x), g_j(x))] \leq [\mathbb{Z}[x] : \mathbf{a}_{j+1}].$$

(The first equality requires $\text{res}(f_j(x), g_j(x)) \neq 0$, which holds because $f_j(x)$ is monic and $g_j(x) \equiv 1 \pmod{p}$ for any $p < 4n^2$.) The coefficient of x^{j-1-i} in $g_j(x)$ is bounded using Hadamard's inequality by $\binom{j-1}{i} (\sqrt{i+1})^{i+1} \|A_{j+1}\|^{i+1} < (j \|A_{j+1}\|)^{i+1}$. Thus all coefficients of $g_j(j \|A_{j+1}\| x)$ are bounded by $(j \|A_{j+1}\|)^j$, the same bound that we used for coefficients of $f'_j(x)$, but with $\|A_{j+1}\|$ in place of $\|A_j\|$ (which does not matter because both are replaced with $t + qr$).

Reproducing (5.8) allows us to deduce (5.10) for $\sum_p \frac{1}{p} v_p([\mathbb{Z}[x] : \mathbf{a}_j])$. Fixing some $p > 4n^2$, we just chose a_j so that $\mathbf{a}_{j-1}$ in (5.9) can be replaced with $\mathbf{a}_j$. This means $-f'_{j-1}(x)$, which is coefficient of b_j in $f'_j(x)$, does not belong to any $\mathfrak{p} | p$ that contains $\mathbf{a}_j$. Thus there is at most one congruence class $b_j \pmod{p}$ per prime $\mathfrak{p} | p$ that violates the condition $(\mathbf{a}_j, f'_j(x), p) = \mathbb{Z}[x]$ implicitly stated in line 10. There are at most $v_p([\mathbb{Z}[x] : \mathbf{a}_j])$ such primes $\mathfrak{p}$. As a result, all of the arithmetic in (5.11) is still correct with $[\mathbb{Z}[x] : \mathbf{a}_j]$ in place of $\text{disc} f_{j-1}(x)$.

The most significant change comes occurs in (5.14) in the effort to avoid those b_j that make $\sum_p \frac{1}{p} v_p(\text{disc} f_j(x))$ too large. Let $\mathbf{b}_j = (f_j(x), f'_j(x))$. Since $|\text{disc} f_j(x)| = |\text{res}(f_j(x), f'_j(x))| = [\mathbb{Z}[x] : \mathbf{b}_j]$, we define $w_{\mathfrak{p}}$ as the minimum of $d(\mathfrak{p}|p) v_{\mathfrak{p}}(\mathbf{b}_j)$ as b_j ranges over $0, q, \dots, \lfloor r \rfloor q$. Also define $w_{\mathfrak{p}}(b)$ as the value of $d(\mathfrak{p}|p) v_{\mathfrak{p}}(\mathbf{b}_j) - w_{\mathfrak{p}}$ when $f_j(x)$ is determined by $b_j = bq$. Unlike the situation with a_j and $\mathbf{a}_j$, both generators of $\mathbf{b}_j$ are affected by b_j . Indeed, $f_j(x) = (x - A_{j,j} - b_j) f_{j-1}(x) + h(x)$ and $f'_j(x) = f'_{j-1}(x) + (x - A_{j,j} - b_j) f'_{j-1}(x) + h'(x)$ for some $h(x)$ of degree at most $j-2$. This shows two things: First, $w_{\mathfrak{p}}$ is bounded by $d(\mathfrak{p}|p) v_{\mathfrak{p}}(\mathbf{b}_{j-1})$, where $\mathbf{b}_{j-1}$

is the ideal generated by the two coefficients of b_j , namely $f_{j-1}(x)$ and $f'_{j-1}(x)$. Therefore,

$$\sum_{\mathfrak{p} \mid p} w_{\mathfrak{p}} \leq v_p([\mathbb{Z}[x] : \mathfrak{b}_{j-1}]) = v_p(\text{disc} f_{j-1}(x)),$$

allowing (5.13) to be replicated. Second, the only primes $\mathfrak{p}$ that can possibly contain $\mathfrak{b}_j$ are those generated by p and an irreducible factor of

$$\begin{aligned} f_{j-1}(x)(f_{j-1}(x) + (x - A_{j,j})f'_{j-1}(x) + h'(x)) - f'_{j-1}(x)((x - A_{j,j})f_{j-1}(x) + h(x)) \\ = f_{j-1}^2(x) + f_{j-1}(x)h'(x) - f'_{j-1}(x)h(x), \end{aligned}$$

which is monic of degree $2j - 2$. Hence the upper bound in (5.14) gets doubled, thereby doubling (5.16). Nevertheless, (5.16) is still less than $\frac{1}{9n}$ for all $n \geq 5$. For $n = 2, 3$, and 4 , (5.15) can be bounded by $\frac{1}{9n}$ by using sharper bounds on $\sum_p \frac{1}{p^2}$ than Lemma 5.8 provides, namely 0.017, 0.0062, and 0.0034, respectively. $\square$

The number of **while** loop iterations in Algorithm 1 is handled similarly:

Lemma 5.10. *With r as defined in Subroutine 3, the **while** loop in Algorithm 1 terminates after at most r iterations.*

Proof. Let $f(x) = f_n(x)$ be the minimal polynomial of λ , let $\mathbf{b}$ assume its value in line 11 of Algorithm 1, and let $b(x) = (1, x, \dots, x^{n-1})\mathbf{b}$. Observe that β is the image of $b(x)$ under $\mathbb{Z}[x] \rightarrow \mathbb{Z}[x]/(f(x)) \xrightarrow{\sim} \mathbb{Z}[\lambda]$. Since the conductor $\mathfrak{f}(\mathbb{Z}[\lambda])$ is contained in the ideal generated by $f'(\lambda)$, we are done if we can show that $(b(x) + a, f(x), f'(x)) = \mathbb{Z}[x]$ for some $a = 0, \dots, [r]$.

For each prime $\mathfrak{p} \mid p$ in $\mathbb{Z}[x]$ that divides $(f(x), f'(x))$, a only needs to avoid at most one congruence class mod p . Every such prime $\mathfrak{p}$ contributes at least one to $v_p(\text{disc} f(x))$, which we know to be zero when $p < 4n^2$ by Lemma 5.4. Replacing $n - 1$ with n in the numerator of (5.7) shows that

$$(5.17) \quad \frac{\log |\text{disc} f(x)|}{r \log n} < \frac{8}{9}$$

when $n \geq 4$. The same bound holds when $n = 2$ or 3 , but we have to be slightly more precise to see it: The lower bound on t is $3q' \log q'$, where $q' \geq n^2 q$. This gives $t > 4.32 \cdot 10^6$ when $n = 2$ and $1.57 \cdot 10^{14}$ when $n = 3$. We use $\binom{n}{i}(t + 2qr)^i$ again to bound the coefficient of x^i in $f(x)$, but now we simply plug that expression directly into the formulas for the discriminant of quadratic and cubic polynomials. The result is $|\text{disc} f(x)| < 8(t + 2qr)^2$ when $n = 2$ and $|\text{disc} f(x)| < 486(t + 2qr)^6$ when $n = 3$. Plugging in the exact values of q , setting $r = n^2 \log t$, and evaluating at our lower bound for t verifies (5.17) for $n = 2$ and 3 .

The exact same chain of inequalities in (5.10), but with (5.17) used in the final step, ends with an upper bound of $\frac{1}{3} + \frac{1}{4}(\frac{8}{9}) = \frac{5}{9}$ instead of $\frac{4}{9}$. Finally, replicating the inequalities in (5.11), but with $\frac{5r}{9}$ used instead of $\frac{4r}{9}$ in the second step, ends with an upper bound of r . Thus there is at least one value of a in the sequence $0, \dots, [r]$ that passes the **while** loop condition. $\square$

Theorem 5.11. *Algorithm 1 with Subroutine 3 runs in deterministic polynomial time.*

Proof. By Lemmas 5.9 and 5.10, all **while** loops terminate in a polynomial number of iterations. Except perhaps line 13 in Algorithm 1 and lines 6 and 10 in Subroutine 3, all steps rely on well-known algorithms that run in polynomial time in their

inputs, which have polynomial bit length in the bit length of M and κ . Methods for line 13 and line 6 and 10 are also established, though perhaps not as well known outside of algorithmic number theory. Ideal or radical ideal membership can be tested in the ring $\mathbb{Z}[x]/(f(x))$ by scaling first computing a $\mathbb{Z}$ -basis for the ideal. (Scale generators over the ring by $1, x, \dots, x^{n-1}$ and reduce the results mod $f(x)$ to produce generators over $\mathbb{Z}$, then reduce to a basis.) Once the ideal becomes a sublattice of $\mathbb{Z}$, we can check the smallest multiple of $1 + 0x + \dots + 0x^{n-1}$ that it contains. If that multiple is 1, line 13 in Algorithm 1 passes the test; if that multiple is a power of 2, line 6 or 10 in Subroutine 3 passes the test. $\square$

Theorem 5.12. *From Algorithm 1 with Subroutine 3, $\beta\alpha_1, \dots, \beta\alpha_n$ is a $\mathbb{Z}$ -basis for an ideal $\beta\mathfrak{a}$ that is coprime to the conductor in the totally real number ring $\mathbb{Z}[\lambda]$. Furthermore, if B is the matrix with i^{th} column $\Sigma(\alpha_i\beta)$ and A is its original value in line 2, then $AB^{-1}\Sigma(\beta\mathfrak{a}) = M\mathbb{Z}^n$, and $\kappa_2(AB^{-1}) < \kappa$.*

Proof. As in the proof of Theorem 4.4, let $A_1 = A$ and B_1 denote the input and output of Subroutine 3, and let $A_2 = B_1$ and $B_2 = B$.

First observe that $\mathbb{Q}(\lambda)$ is a totally real number field of degree n . Indeed, all perturbations found in the **for** loop of Subroutine 3 are even and preserve symmetry. Thus $B_1 \equiv S_2 \pmod{2}$ by line 2 of Subroutine 3. From Lemma 5.1, the characteristic polynomial of S_2 is irreducible mod 2. Hence the characteristic polynomial of $A_2 = B_1$, which is the minimal polynomial of λ is irreducible with only real roots.

Next, coprimality of β and $\mathfrak{f}(\mathbb{Z}[\lambda])$ is ensured by the **while** loop in Algorithm 1. Line 10 in Subroutine 3 shows that $(\mathfrak{a}, f'(\lambda))$ is not contained in any prime ideals above some $p > 4n^2$, which implies the same for $(\mathfrak{a}, \mathfrak{f}(\mathbb{Z}[\lambda]))$. For primes $p < 4n^2$, we have Lemma 5.4.

Finally, we check that $\kappa_2(A_i B_i^{-1}) < \sqrt{\kappa}$ for $i = 1, 2$. The proof for $A_2 B_2^{-1}$ is almost identical to that in the proof of Theorem 4.4. Let us use the same notation for the diagonal matrices C and C' . The only difference now is that the initial upper bound of 1 in line 10 of Algorithm 1 can grow throughout the **while** loop in lines 12–13, but it is capped at $\lceil r \rceil$ by Lemma 5.10. This is perfectly accounted for by the presence of r on the right side of line 7 in Algorithm 1. So the analysis in (4.2) holds, and we obtain $\kappa_2(A_2 B_2^{-1}) < \sqrt{\kappa}$.

To deal with $A_1 B_1^{-1}$, we need an analog of (4.1). Let t denote the value of “ $\|A\|_2$ ” in line 3 of Subroutine 3, just as in the proof of Lemma 5.9. By line 2,

$$t \geq 3q'\|A_1\|_2 \log(q'\|A_1\|_2) - \|S_q\| \geq 3q'\|A_1\|_2 \log(q'\|A_1\|_2) - \frac{q'}{2n},$$

which implies

$$(5.18) \quad \frac{t}{\log t} > \frac{5}{2}q'\|A_1\|_2.$$

The difference between $3q'\lceil \log(q'\|A_1\|_2) \rceil A_1$ and B_1 consists of entries from S_q and a_j and b_j . Lemma 5.9 bounds a_j and b_j by qr , and the entries of S_q are bounded in magnitude by $\frac{q}{2}$. There are at most two perturbations a_j and b_j in each row and column of B_1 . Since the 1-norm and ∞ -norm of a matrix are bounded by the maximum sum of entry magnitudes among its columns and rows, we have

$$\left\| A_1 - \frac{B_1}{3q'\lceil \log(q'\|A_1\|_2) \rceil} \right\|_2 \leq \max_{p=1, \infty} \left\| A_1 - \frac{B_1}{3q'\lceil \log(q'\|A_1\|_2) \rceil} \right\|_p$$

$$\begin{aligned}
&\leq \frac{2qr + nq/2}{3q' \lceil \log(q' \|A\|_2) \rceil} \\
&< \frac{5\|A_1\|_2 qr}{2(3q' \|A_1\|_2 \lceil \log(q' \|A_1\|_2) \rceil + nq/2)} \\
&\leq \frac{5\|A_1\|_2 q n^2 \log t}{2t} && \text{by definition of } r \text{ and } t \\
&< \frac{qn^2}{q'} && \text{by (5.18)} \\
&\leq \frac{\sqrt{\kappa} - 1}{\|A_1^{-1}\|_2(\sqrt{\kappa} + 1)} && \text{by definition of } q'.
\end{aligned}$$

It follows that $\kappa_2(A_1 B_1^{-1}) < \sqrt{\kappa}$ by Lemma 3.6. $\square$

This justifies Theorem 1.1 in the introduction. We can also now prove the claimed hardness result from the introduction.

Corollary 5.13. *Restricted to the canonical embedding of ideals coprime to the conductor in totally real, monogenic number rings, SVP and CVP are NP-hard to approximate within factors of $\sqrt{2}$ and $n^{c/\log \log n}$, respectively.*

Proof. This is almost immediate from Theorems 5.11 and 5.12. It remains only to check that the value of κ needed to preserve the factors $\sqrt{2}$ and $n^{c/\log \log n}$ is not so small that it destroys running time.

In [Wan26], Wan proves that $\sqrt{2}$ -SVP is NP-hard in the setting of integer lattices. In other words, Proposition 3.5 applies—we may perfectly preserve the factor $\sqrt{2}$ by making κ sufficiently small in Algorithm 1. Note that the bound Proposition 3.5 places on κ gives $\lceil (\sqrt{\kappa} - 1)/(\sqrt{\kappa} - 1) \rceil$ bit length that is polynomial in the bit length of A from line 2 of Algorithm 1, which is polynomial in the bit length of M using Storjohann’s Smith Normal Form algorithm [Sto00].

We need not bother with Proposition 3.5 to prove NP-hardness of CVP. Dinur’s factor of $n^{c/\log \log n}$ is not fixed like $\sqrt{2}$. For example, $\kappa = 2$ works. $\square$

REFERENCES

- [AD97] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/ average-case equivalence. *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*. 1997, pp. 284–293. DOI: [10.1145/258533.258604](https://doi.org/10.1145/258533.258604).
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems. *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. 1996, pp. 99–108. DOI: [10.1145/237814.237838](https://doi.org/10.1145/237814.237838).
- [Ajt98] Miklós Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions. *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*. ACM, 1998, pp. 10–19. DOI: [10.1145/276698.276705](https://doi.org/10.1145/276698.276705).
- [Bau+17] Jens Bauch, Daniel J. Bernstein, Henry de Valence, Tanja Lange, and Christine van Vredendaal. Short generators without quantum computers: The case of multiquadratics. *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2017, pp. 27–59. DOI: [10.1007/978-3-319-56620-7_2](https://doi.org/10.1007/978-3-319-56620-7_2).
- [Ben23] Huck Bennett. The complexity of the shortest vector problem. *ACM SIGACT News* 54.1 (2023), pp. 37–61.
- [Bha14] Manjul Bhargava. The geometric sieve and the density of squarefree values of invariant polynomials. [arXiv:1402.0031](https://arxiv.org/abs/1402.0031). 2014.

- [Bol+19] Madalina Bolboceanu, Zvika Brakerski, Renen Perlman, and Devika Sharma. Order-LWE and the hardness of ring-LWE with entropic secrets. *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2019, pp. 91–120.
- [BP22] Huck Bennett and Chris Peikert. Hardness of the (Approximate) Shortest Vector Problem: A Simple Proof via Reed-Solomon Codes. *International Workshop on Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques*. 2022. URL: <https://api.semanticscholar.org/CorpusID:246867325>.
- [BPR06] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. Berlin, Heidelberg: Springer, 2006. doi: [10.1007/3-540-33099-2](https://doi.org/10.1007/3-540-33099-2).
- [BSW22] Manjul Bhargava, Arul Shankar, and Xiaoheng Wang. Squarefree values of polynomial discriminants I. *Inventiones mathematicae* 228 (2022), pp. 1037–1073. doi: [10.1007/s00222-022-01098-w](https://doi.org/10.1007/s00222-022-01098-w).
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient Fully Homomorphic Encryption from (Standard) LWE. *Proceedings of the IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2011, pp. 97–106. doi: [10.1109/FOCS.2011.12](https://doi.org/10.1109/FOCS.2011.12).
- [CDW21] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time. *Journal of the ACM (JACM)* 68.2 (2021), pp. 1–26.
- [CGS14] Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale. docbox.etsi.org/Workshop/2014/201410_CRYPT0/S07_Systems_and_Attacks/S07_Groves_Annex.pdf. Presented at the ETSI 2nd Quantum-Safe Crypto Workshop, 2014.
- [Cra+16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2016, pp. 559–585. doi: [10.1007/978-3-662-49890-3_21](https://doi.org/10.1007/978-3-662-49890-3_21).
- [Dic01] Leonard Eugene Dickson. *Linear Groups: With an Exposition of the Galois Field Theory*. Leipzig: B. G. Teubner, 1901.
- [Din02] Irit Dinur. Approximating CVP to within almost-polynomial factors is NP-hard. *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*. ACM, 2002, pp. 266–275. doi: [10.1145/509907.509950](https://doi.org/10.1145/509907.509950).
- [Dus18] Pierre Dusart. Explicit estimates of some functions over primes. *The Ramanujan Journal* 45.1 (2018), pp. 227–251.
- [Ebe22] Sean Eberhard. The characteristic polynomial of a random matrix. *Combinatorica* 42.4 (2022), pp. 491–527.
- [Eke91] Torsten Ekedahl. An infinite version of the Chinese remainder theorem. *Commentarii mathematici Universitatis Sancti Pauli= Rikkyo Daigaku sugaku zasshi* 40.1 (1991), pp. 53–59.
- [Emd81] Peter van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Report 81-04. Mathematical Institute, University of Amsterdam, 1981.
- [Fer+23] Asaf Ferber, Vishesh Jain, Ashwin Sah, and Mehtaab Sawhney. Random symmetric matrices: rank distribution and irreducibility of the characteristic polynomial. *Mathematical Proceedings of the Cambridge Philosophical Society* 174.2 (2023), pp. 257–270.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing* (2009), pp. 169–178. doi: [10.1145/1536414.1536440](https://doi.org/10.1145/1536414.1536440).

- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. *International Algorithmic Number Theory Symposium (ANTS-III)*. Vol. 1423. Lecture Notes in Computer Science. Springer, 1998, pp. 267–288. DOI: [10.1007/BFb0054868](https://doi.org/10.1007/BFb0054868).
- [HS26] Isaac Hair and Amit Sahai. SVP_p is NP-hard for all $p > 2$, Even to approximate Within a factor of $2^{\log^{1-\epsilon} n}$. [arXiv:2511.04125](https://arxiv.org/abs/2511.04125). 2026.
- [Len92] Hendrik W Lenstra. Algorithms in algebraic number theory. *Bulletin of the American Mathematical Society* 26.2 (1992), pp. 211–244.
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. *Lecture Notes in Computer Science* (2006), pp. 144–155. DOI: [10.1007/11787006_13](https://doi.org/10.1007/11787006_13).
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2010, pp. 1–23.
- [mat20] GreginGre (math.stackexchange.com/users/447764/gregingre). Existence of orthogonal base for finite Galois extension over characteristic 2. Mathematics Stack Exchange. math.stackexchange.com/questions/3534461 (accessed: 2026-09-02). 2020.
- [MG12] Daniele Micciancio and Shafi Goldwasser. Complexity of lattice problems: a cryptographic perspective. Vol. 671. Springer Science & Business Media, 2012.
- [Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity* 16 (2007), pp. 365–411. DOI: [10.1007/s00037-007-0234-9](https://doi.org/10.1007/s00037-007-0234-9).
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem. *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*. 2009. DOI: [10.1145/1536414.1536461](https://doi.org/10.1145/1536414.1536461).
- [Pei16] Chris Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science* 10.4 (2016), pp. 283–424.
- [PHS19] Alice Pellet-Mary, Guillaume Hanrot, and Damien Stehlé. Approx-SVP in ideal lattices with pre-processing. *Lecture Notes in Computer Science* (2019), pp. 685–716. DOI: [10.1007/978-3-030-17656-3_24](https://doi.org/10.1007/978-3-030-17656-3_24).
- [PP24] Chris Peikert and Zachary Pepin. Algebraically Structured LWE, Revisited: C. Peikert, Z. Pepin. *Journal of Cryptology* 37.3 (2024), p. 28.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. *Theory of Cryptography Conference*. Springer. 2006, pp. 145–166.
- [PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 461–473.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*. ACM, 2005. DOI: [10.1145/1060590.1060603](https://doi.org/10.1145/1060590.1060603).
- [RS62] J. Barkley Rosser and Lowell Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois Journal of Mathematics* 6 (1962). DOI: [10.1215/ijm/1255631807](https://doi.org/10.1215/ijm/1255631807).
- [RSW18] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-LWE and polynomial-LWE problems. *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2018, pp. 146–173.
- [Sho90] Victor Shoup. New algorithms for finding irreducible polynomials over finite fields. *Mathematics of Computation* 54.189 (1990), pp. 435–447.

- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing* 26 (1997), pp. 1484–1509. DOI: [10.1137/S0097539795293172](https://doi.org/10.1137/S0097539795293172).
- [Ste+09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. *International Conference on the Theory and Application of Cryptology and Information Security*. Springer. 2009, pp. 617–635.
- [Sto00] Arne Storjohann. Algorithms for matrix canonical forms. PhD thesis. ETH Zurich, 2000.
- [Wae34] Bartel Leendert van der Waerden. Die Seltenheit der Gleichungen mit Affekt. *Mathematische Annalen* 109.1 (1934), pp. 13–16.
- [Wan26] Daqing Wan. NP-hardness of SVP in Euclidean space. [arXiv:2603.27398](https://arxiv.org/abs/2603.27398). 2026.
- [WY16] Wei Wang and Tao Yu. Square-free discriminants of matrices and the generalized spectral characterizations of graphs. [arXiv:1608.01144](https://arxiv.org/abs/1608.01144). 2016.

Email address: `dem6@clemson.edu`

CLEMSON UNIVERSITY, O-110 MARTIN HALL, 2020 PARKWAY DRIVE, CLEMSON, SC